

The Best Answers? Think Twice: Online Detection of Commercial Campaigns in the CQA Forums

Cheng Chen, Kui Wu, S. Venkatesh
Department of Computer Science
University of Victoria
Victoria, Canada

Kesav Bharadwaj R
Department of Computer Science
Bits-Pilani
Pilani, India

ABSTRACT

In an emerging trend, more and more Internet users search for information from Community Question and Answer (CQA) websites, as interactive communication in such websites provides users with a rare feeling of trust. More often than not, end users look for instant help when they browse the CQA websites for the best answers. Hence, it is imperative that they should be warned of any potential commercial campaigns hidden behind the answers. However, existing research focuses more on the quality of answers and does not meet the above need. In this paper, we develop a system that automatically analyzes the hidden patterns of commercial spam and raises alarms instantaneously to end users whenever a potential commercial campaign is detected. Our detection method integrates semantic analysis and posters' track records and utilizes the special features of CQA websites largely different from those in other types of forums such as microblogs or news reports. Our system is adaptive and accommodates new evidence uncovered by the detection algorithms over time. Validated with real-world trace data from a popular Chinese CQA website over a period of three months, our system shows great potential towards adaptive online detection of CQA spams.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: General; J.4 [Social and Behavioral Science]: Sociology

General Terms

Design, Experimentation, Measurement

Keywords

CQA forums; Online detection; Paid posters

1. INTRODUCTION

Web 2.0 social websites are playing an increasingly important role on the Internet by utilizing the wisdom of crowds [24].

One such example is the Community Question and Answer (CQA) portals on which users can post and answer questions, such as Yahoo! Answers, Naver and Baidu Zhidao [4, 30, 20]. Some CQA websites like Quora [23] attract users by offering professional answers, most of which come from verified people in reality. These websites gain popularity and trust by providing a sense of interaction between the questioner and the masses. With millions of archived Q&A sessions [29], CQA forums have become a major source of advice for many Internet users.

As a large knowledge base of crowds, the archived Q&A sessions have been used for automatic question answering and recommendation. Nevertheless, the quality of user-generated content in the Q&A sessions varies drastically. For instance, some answers do not match the questions and even contain spam and rude words. In recent years, tremendous efforts have been made to locate better answers and remove spam from the archived questions and answers resource. Techniques such as analysis of text, user-question-answer's link relationship, and user feedback features have been used in tools like PageRank to identify high-quality web pages [13, 15, 2].

Existing techniques, however, may not work well in the presence of the so-called Internet water army, a large crowd of hidden posters who get paid to generate artificial content in the social media for commercial profits. Paid posters have become popular with the booming of crowd-sourcing marketing. As confirmed in [26], crowd-sourcing systems such as Amazon's Mechanical Turk, Zhu Ba Jie (a similar Chinese crowd-sourcing site), have been broadly used for commercial campaigns. Due to their popularity, the CQA forums have become the targets of those campaigns that create untruthful Q&A sessions for commercial purpose. Consider the following example:

Question: I tried several methods to lose weight but all failed. What should I do? Please give me some advice!

Best answer: Don't worry, I have experienced the same pain as you. Firstly, you have to keep a healthy diet. Be careful about the nutrition in your food and never eat fast food. Secondly, don't sit too long in front of a computer. Finally, perform physical exercise everyday. What's more, you can also try a product named X. This product contains ingredients such as ... and can help you lose weight without any risks.

The above Q&A session was actually generated by paid posters. The answer provides very practical advice at first and then gives suggestion on the product which needs to be promoted. The practical advice part is to earn the trust of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

the users. We have observed that fake answers generated by paid posters are often long enough and quite relevant to the questions, and some paid posters involved in the fake Q&A sessions are ranked high according to the website's reputation system.

Based on textual similarities, previous work [18, 6, 7] is likely to treat the above answer as of high quality due to the high relevance of textual features between the answer and question content. As a result, the output may contain commercial spam, resulting in a credibility problem. Therefore, additional strategies, such as *writing templates*, *public calls for commercial campaigns*, and *a poster's track reputation*, should be integrated for the effective detection of paid posters. Furthermore, most existing work relies on offline analysis, while end users demand for instant help and should be warned of potential commercial campaigns when they browse a CQA forum. The call for a real-time response system that can detect potentially fake Q&A sessions on the fly is strong.

We tackle the above two challenges in this paper by designing an adaptive online detection system tailored specifically for CQA forums. Our contributions are as follows:

- We discover that the behavioral features of paid posters are different in CQA forums when compared to other types of forums such as microblog (also called Weibo, a Twitter like service in China) and news reports. We identify the special features of paid posters in CQA forums that are useful in the detection.
- Based on the identified special features, we design a detection method which uses machine-learning techniques and assigns credibility scores to each of the best answers by using semantic analysis and user features, such as users' history data.
- We implement an adaptive, online detection system which automatically analyzes the hidden patterns of commercial spams and raises alarms instantaneously to end users whenever a potential commercial campaign is detected. Our system is adaptive and accommodates new evidence gathered by the detection algorithms over time.

2. DATA COLLECTION AND LABELING

2.1 How Do Online Paid Posters Work in CQA Portals

To understand the background, we start with a brief introduction on how online paid posters work in CQA sites.

With the advent of popular crowd-sourcing websites, companies tend to hire paid posters to help them hype their products in the social media. Research [26] has shown that paid posters are capable of generating large information cascades that could escape security check and accelerate spam dissemination on social media, like microblogging services and community-based question and answer websites.

Our research is based on Baidu Zhidao, a Chinese CQA website that is similar to Yahoo! Answers. During our study on the CQA-oriented promoting campaigns on crowd-sourcing websites, we discovered specification with detailed requirements and templates for the paid posters. The requirements provide not only basic description regarding the product but also types of paid posters needed. For example,

some companies request that posters should have a good reputation. Note that many CQA websites have a reputation system and assign high-level reputation indicators to "trustworthy" users whose answers are mostly selected as the best answers. Those reputation systems track the history of users but are not designed to analyse and detect online paid posters.

It is very interesting to notice that companies that hire paid posters also provide several templates for questions and answers. For instance, in a medicine promotion case, the question describes a certain symptom, and the answer explains reasons for the symptom and recommends taking the specific medicine. Both question and answer templates are carefully crafted to sound real. The answers usually include personal experience with the products. In addition, the instructions will advise paid posters to insert their own sentences in the templates rather than just copying and pasting the templates.

Using these templates, paid posters can create complete Q&A sessions. They first pose a question, and use a different user ID to post the answer. This could be achieved by one user registering for multiple IDs or by several colluding posters. They then select the answer as the best answer, after waiting for other users to post answers. This waiting time is to cheat the readers into believing that the best answer is chosen from many answers. In the CQA portals, once the best answer is decided, the Q&A session is considered *closed* and no new answers can be added to the session.

2.2 Data Collection

Users who register on Baidu Zhidao participate in various Q&A sessions, either as question askers or repliers. Since we already know that paid posters who accept missions from crowd-sourcing sites create a variety of Q&A sessions on the site for product propaganda, the collecting process can be targeted directly to the product campaigns. In addition, since the readers tend to pay more attention to the best answers and also due to the manner in which online paid posters are supposed to work, we only collected the best answers and ignored other ones. This is to avoid collecting a large amount of irrelevant information for this study.

In order to collect campaign Q&A sessions, we first visited the crowd-sourcing websites, where the paid posters apply for campaign tasks and get paid, as stated in Section 1. After going through campaigns calling for paid posters, we selected 11 closed requests because the paid posters who worked for the 11 products had finished the tasks. We extracted keywords from the 11 products and searched for Q&A sessions with the keywords on Baidu Zhidao. We used a crawler to collect all the links from the searching result. We then used another crawler to visit and download the web pages associated with the links. The results included not only the campaign sessions, but also normal sessions containing the keywords. After parsing all the collected web pages, we obtained a group of target users, including both paid posters and normal users, as well as the links to the users' homepages hosted by Baidu Zhidao.

By following the users' homepages, we could find useful information for our research. For example, a user's homepage has the question-answer history of this user, and includes all the Q&A sessions where this user posted his/her answers. The question-answer history provides a good knowledge on the multiple campaigns that a potential paid poster might

have been involved.

Having obtained the initial dataset of IDs and links, we then visited each user’s homepage, retrieved every Q&A session that the user participated in. We only collected the closed Q&A sessions (i.e., the best answer determined). A closed Q&A session implies that users can no longer post new answers to the question, but they can click the “Like” button to support the posted answers, including the best answer and other answers. From those Q&A sessions, we finally extracted information used in our analysis. The recorded information from those web pages includes *questioner ID*, *answer ID*, *time*, *title*, *question content*, *answer content*, *user feedbacks* (*visited times*, *ratings*).

From the Q&A website, *Baidu Zhidao*, we crawled, 6462 users’ question-answer history records accumulated during a three-month period from October to December in 2011. For each user, we built a list of history information, showing the question, answer, participated user IDs, and other features. Associated with the 6462 user IDs, we have 75,200 Q&A sessions in total, all having the best answer.

2.3 Manual Data Labeling

To get a sample dataset for feature analysis, campaign sessions should be differentiated from the normal ones. By reading the best answers, we manually labeled the Q&A sessions in the dataset. The labeling process mainly depends on the Q&A templates from the crowd-sourcing websites such as Zhubajie [31] and Tiancaicheng [25]. We summarize the applied techniques below:

1. Since we have collected a list of 11 products which were hyped in the Baidu Zhidao, we could compare the Q&A content with the campaign templates. If the product’s name is in the 11 initial samples and the contents match the templates, such as the descriptive words and the organized pattern of sentences, we labeled it as a campaign Q&A session. We stress that there is difference between our work and related research which needs to judge the quality of answers. The evaluation of quality of answers is usually based on question-answer relevance, length of the texts, grammar correctness, politeness, and so on. To obtain a reliable dataset, researchers often rely on multiple assessors and are faced with the difficulty of reaching an agreement among the multiple evaluation results. Our labeling method differs from the above and largely avoids the annotation difficulty, because we know exactly the name of the hyped product and how paid posters would write the Q&A sessions.
2. When we encountered new products not in the list of 11 initial samples, we recorded the product’s name and searched it in the crowd-sourcing websites. If we found the template of this product, we use the above method to compare their contents.
3. If a new product is listed in the campaign websites but the template is not available, we followed some special features normally found in Email spam to make a decision. For example, a spam may use different fonts to write the telephone numbers and insert special characters between the product’s name. This type of operations is usually used to escape detection by the filter system. We labeled the session as campaign if

the product’s name is in a campaign list and the best answer has special features similar to Email spam.

4. If we could not find the new product in the campaign websites, we then tried to identify potential templates used in the same category of products and special features obvious in an Email spam. If none of those could be identified, we labeled the session as a normal session.

Up to now, we have labeled 4998 samples in our dataset. Among these, 2147 samples are campaign Q&A sessions and the other 2851 samples are normal ones. The sample size is large enough for our current study. Since we selected 11 campaigns, which were posted on the crowdsourcing websites, as the seeds of our crawler, the proportion of campaign sessions is relatively high in the dataset.

When we manually labeled our datasets, we carefully read the contents of a user’s post. The meaning can be understood by human but is hard to use in machine learning based classification. Even with the above template based labeling method, it is not easy to write an algorithm to automatically identify a campaign session because a poster may re-phrase the template in their own words. Due to these reasons, we need to search for statistical features that can be effectively used towards building a detection system.

3. ANALYSIS OF STATISTICAL FEATURES

3.1 Insufficiency of Existing Statistical Features

We firstly demonstrate the difficulty of the problem we study by analyzing existing features, some of which have been used in related research such as evaluation of high quality answers or detection of Internet water army in news report websites [8] and showing their limitations.

3.1.1 Interval Post Time

In [19], Arjun *et al.* defined several spamming indicators for modelling the behaviour of fake review writers. They found that spammers of a spam group tend to post reviews during a short time interval. This feature has been shown to be a good indicator to detect Internet water army in news report websites [8].

In our work, we consider two time stamps for a Q&A session: One is the time when the questioner post the question topic (the ask time), and the other one is the time when the best answer is posted by a replier (the best answer posted time). We define *interval post time* as the latter time stamp minus the former one.

In Figure 1, we show the probability distribution of interval post time with “pp” (abbreviation of paid posters) for campaign sessions and “nu” (abbreviation of normal users) for non-campaign sessions. The x-axis is drawn by *lg* scale.

From the figure, we find it difficult to tell the difference between campaign and non-campaign Q&A sessions. Two reasons may contribute to the above phenomenon. There are many normal users who spend much time on the Q&A website and try to post answers to *open* questions, especially those questions associated with some *rewards points*. These people are known as *bounty hunters*. Most bounty hunters post very good answers because they want to get more rewards points. On the other hand, online paid posters, before they post and choose the best answer, normally wait

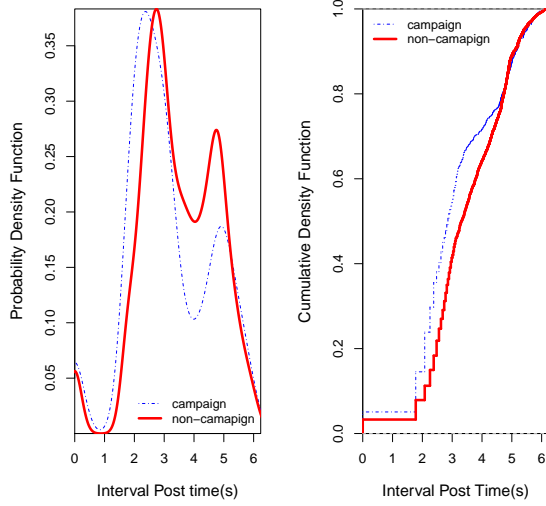


Figure 1: The PDF and CDF of the interval post time

for some random time for other answers appearing in the session. This is to give readers a fake impression that the best answer is selected among many answers. While paid posters try to finish a job as quickly as possible in news review websites [8], the same behaviour does not exist here.

3.1.2 Number of Other Answers

Before the question is closed, users can post their own answers. This variable counts the number of answers other than the best one. Intuitively, if the paid posters create the sessions themselves, they may not have patience to wait for more replies. They could close the sessions and get paid as soon as possible. To test this conjecture, we show the probability distribution of this feature for campaign sessions and normal sessions in Figure 2.

Similar to the interval post time, the number of other answers does not indicate much difference for the two types of Q&A sessions. This invalidates the above conjecture and we do not consider it a good feature for the detection of paid posters in CQA portals.

3.1.3 Number of Likes

Similar to the “Like” button in Facebook, if other readers find the best answer to be helpful, they may click the “like” button. The number on the button indicates the total number of clicks. Intuitively, this feature represents user’s feedback and should be helpful in identifying trustful answers. The more “likes” an answer receives, the more likely it is a good answer. However, as shown in Figure 3, this is not a reliable feature. This is because the paid posters could click the button themselves and even use different user IDs to click multiple times. This behavior is also confirmed in [5] as the “vote spam attack”.

3.1.4 Relevance between Questions and The Best Answers

This feature is extensively used before in identifying high-quality answers [18, 6, 2, 7, 22]. The previous work is usually

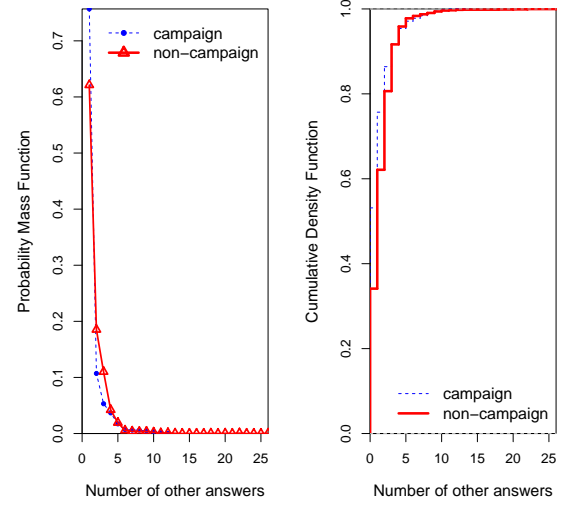


Figure 2: The PMF and CDF of the number of other answers




Figure 3: The PMF and CDF of the number of likes

based on following assumptions:

1. Semantically high relevance between questions and answers indicates high quality.
2. Selected best answers should have higher quality than other answers.

The above assumptions are risky for the detection of potential campaigns created by paid posters. In commercial campaigns, answers with *high-quality* are rather misleading and would beat the retrieval mechanism. Many of the answers are well-organized and highly related to the questions. In this sense, a “high-quality” answer does not necessarily mean trustworthiness. Thus, we do not consider the relevance measure in our work.

3.2 Special Features for CQA Portals

The limitations of existing statistical features shown above led us to look for new features specific to users in CQA websites.

3.2.1 Spam Grade of Questioner ID (SGqID)

It indicates whether the questioner tends to ask campaign questions. For a given questioner ID (qID), we calculate the ratio of the number of campaign sessions and the total number of sessions in which the user has participated,

$$SGqID = \frac{q_1}{q_0 + q_1} \quad (1)$$

where q_0 and q_1 are the number of non-campaign and campaign sessions where the user appears as the questioner, respectively. If the system does not record such information, we set its SGqID value to 0.5.¹

3.2.2 Spam Grade of Answerer ID (SGaID)

It indicates whether the best answer poster tends to write campaign answers. For a given answerer ID (aID), we calculate the ratio of the number of campaign sessions and the total number of sessions in which the user has participated,

$$SGaID = \frac{a_1}{a_0 + a_1} \quad (2)$$

where a_0 and a_1 are the number of non-campaign and campaign sessions the user appears as the poster of the best answers, respectively. Similar to SGqID, if the system does not record such information, we set its SGaID value to 0.5.

3.2.3 Spam Grade of the Text (SGtext)

It indicates whether the collection of words in sessions associated to a user tends to be campaign specific. To calculate this feature, we need to perform statistical analysis over the words. Text information of a Q&A session consists of the title, the content of question, and the content of the best answer. We remove the duplicate words so that we can get a collection of distinct words, $word_1, word_2, word_3 \dots word_n$, for each Q&A session. For each word, we calculate *spam grade* which characterizes the property of the word, i.e., whether it is more campaign oriented or non-campaign oriented. Words with higher benchmark are more likely to imply hidden promotion behavior. To get rid of the impact of different length, we take the average value over the summation of the benchmarks of all words as the spam grade of the whole text. For each word, the definition of spam grade goes like this:

$$SGword_i = \lg \left(\frac{N+1}{n_i+1} \right) * \frac{s_i+1}{S+1} \quad (3)$$

where N and S are the total number of non-campaign and campaign sessions in the databases and n_i and s_i are the number of non-campaign and campaign sessions where the $word_i$ appears. The term “+1” is used to normalize the result in case of zero counts. Then the spam grade of text with L distinct words is calculated as:

$$SGtext = \frac{SGword_1 + SGword_2 + \dots + SGword_L}{L} \quad (4)$$

¹This decision follows the Maximum Entropy Principle [16], i.e., we should “make use of all the information that is given and scrupulously avoid making assumptions about information that is not available.”

3.3 Property of the Feature Set

Figure 4 exhibits the values using the three “SG” features in the previous section.




Figure 4: 4998 samples captured by SGqID, SGaID and SGtext

Through this figure, we can observe a clear gap between the campaign sessions and non-campaign sessions. We can then apply regression based techniques to calculate the campaign score, which indicates whether a Q&A session tends to be a campaign.

4. DETECTION METHOD

In this section, we introduce a logistic regression approach [10, 1] to calculate campaign scores for Q&A sessions using the three proposed “SG” features.

4.1 The Algorithm

Figure 4 has already shown that the samples can be distinguished by the three proposed features, SGqID, SGaID and SGtext. In order to get a score indicating whether a Q&A session is a potential commercial campaign or not, we apply logistic regression as the learning method. We can use it to calculate values of $P(Y = 1|X, \theta)$ and $P(Y = 0|X, \theta)$. Here, Y is a indicator variable, where $Y = 1$ and $Y = 0$ represent campaign and non-campaign Q&A sessions, respectively. \mathbf{X} is a vector of three features for each session. θ is a vector of model parameters, each associated with a session feature and including an individually constant item which is not related to the session features.

By applying the sigmoid function, the hypothesis $h_{\theta}(\mathbf{X})$ which outputs a score of $P(Y = 1|X, \theta)$ or $P(Y = 0|X, \theta)$ (termed as *campaign score*) is defined as follows:

$$h_{\theta}(\mathbf{X}) = \frac{1}{1 + e^{-\theta^T \mathbf{X}}} \quad (5)$$

where $\theta^T \mathbf{X} = \theta_1 + \theta_2 * SGqID + \theta_3 * SGaID + \theta_4 * SGtext$. To facilitate the matrix calculation, we add an all-1 column to \mathbf{X} .

In practice, the higher the score, the higher the probability that the given session is a campaign session. The values of θ will be learned by logistic regression. The objective then becomes a regression problem where we optimize the model so that the output campaign scores of sessions are close to their true labels (0 or 1).

The convex cost function of this optimization problem is given by

$$J(\theta) = \frac{1}{m} \sum_{i=1}^m [-y^{(i)} \log(h_{\theta}(x^{(i)})) - (1-y^{(i)}) \log(1-h_{\theta}(x^{(i)}))] \quad (6)$$

where m is the number of samples in the training dataset and x is a matrix consisting of m feature vectors of the training samples. We use gradient descent method to find the minimum of the cost function and the corresponding values in θ .

4.2 Regression and Classification Results

We shuffled the 4998 labelled samples and took 3500 of them as training set and the remaining 1498 as test set. The statistics of the two datasets are shown in Table 1. Note that the split of the dataset is arbitrary only to illustrate our detection method and results. We will demonstrate how our system adapts with data changes for real-time detection in the next section.

Table 1: The number of samples of training and test datasets

	non-campaign	campaign
training	1984	1516
test	867	631

When the θ is optimized, we then calculate the campaign score of each Q&A session in the test dataset. The distribution of scores for normal sessions and campaign sessions is shown in Figure 5.




Figure 5: Scores of test set (CDF)

From the figure, we can see that the two types of Q&A sessions exhibit great difference on the distribution of the campaign scores. Most of the campaign scores are very close to their true labels ($Y = 1$ or $Y = 0$). Using the scores, we can either provide the raw scores to the users to help them make decisions when reading the answers, or we can assign the labels based on a threshold value, i.e., $Y = 1$ when the campaign score is larger than the threshold value and $Y = 0$ otherwise.




Figure 6: ROC Curve of the classification result with different threshold values

In Figure 6, we show the ROC curve based on different threshold values, 0.1, 0.2, ..., 0.9. The points on the curve are mostly located at the top left position of the curve. The reason is that the campaign scores of most campaign sessions are higher than 0.9 while the campaign scores of most normal sessions are smaller than 0.1. This curve shows that the system performance is robust with a large range of threshold value.

Based on Figure 5 and Figure 6, we set 0.5 as our current threshold. With this threshold, we evaluate the following four performance metrics:

$$\begin{aligned} \text{Precision} &= \frac{\text{TruePositive}}{\text{TruePositive} + \text{FalsePositive}} \\ \text{Recall} &= \frac{\text{TruePositive}}{\text{TruePositive} + \text{FalseNegative}} \\ F\text{-measure} &= 2 * \frac{\text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} \\ \text{Accuracy} &= \frac{\text{TrueNegative} + \text{TruePositive}}{\text{TotalNumberOfUsers}} \end{aligned}$$

The classification results are shown in Table 2. Based on the performance results in the table, we can see that the proposed features for detecting campaign sessions are valid and effective. Though the performance of our offline analysis is very satisfactory, we will test its performance on other and bigger datasets for further validations of our feature set in the future. In the next section, we will introduce an adaptive online detection system which adaptively learns data changes over time and return detection results in real time.

Table 2: Performance Results

Precision	Recall	F-measure	Accuracy
98.90%	99.68%	99.29%	99.40%

5. ADAPTIVE ONLINE DETECTION SYSTEM

In the previous section, we have shown that we can build a model to effectively calculate the campaign score and predict the labels of unknown sessions. In practice, however,

this offline analysis does not work well for users who would like to be advised of potential campaigns in real time. This requirement encourages us to design an online version of detection system, which can return campaign scores and/or predicted results in real time. We therefore build a prototype of such an adaptive online detection system. The word “adaptive” implies that this system can update its database using new samples and generate new model parameters.

5.1 Overview of System Design

The major components of the detection system include browser plugin and a remote server. Figure 7 shows the system architecture and the communication between the client plugin and the server.

As shown in Figure 7, the sequence of actions that take place when a user opens a Q/A session are:

1. The plugin first sends only the URL of the page to the server. The server searches for the url in its database. If it is found, the server returns the score (spam rating) to the client. The client side script displays the result. This avoids unnecessarily sending complete web page to the server if it is already present in the database.
2. If the URL is not present, the server sends a response *not found* and the client after receiving the response sends the rest of the data to the server through another *XMLHTTPRequest* and waits for the server’s response.
3. The server receives the data, segments the text into words, and stores it in the database. The server then extracts the statistical features necessary for the analysis from the data. Logistic regression analysis is performed to predict the class of the session (spam or no spam). If the session is classified as a spam, an alert is returned back to the user.
4. The client-side script displays the result to the user.
5. (Optional) If the user is an authorized user, the user can provide feedback to the server (whether or not he/she feels the session is a campaign session). There are three types of users in the system: regular users are those who use our system and they are not granted the right to annotate sessions; helper users are those who have experience and are capable of helping label the data; the administrator is the person responsible for the management of the system. Note that helpers could be contracted out to employees of professional companies such as Rediff Shopping and eBay [19].
6. When newly labelled sessions are available, the system updates the detection model using existing and newly labelled data. Note that this step could be done regularly in a daily or even weekly basis.

5.2 Plugin Design

The plugin is a Google Chrome extension. It must be installed on the Chrome browser in the user’s system. The plugin consists of *manifest.json* file, a HTML file and a *contentscript.js* file. The *contentscript.js* file specifies the javascript to be executed on the webpage the user is browsing. The *manifest.json* file contains information regarding the name, version of plugin and the HTML, script files associated with the plugin. The manifest file also contains a list of permissions that the plugin might use to access servers.

The functions of the plugin can be separated into three major steps.

1. Extract data from the webpage. All the data required from the webpage are extracted from the HTML source of the webpage. Separate javascript functions were written for extracting various information. The information extracted includes the page URL, Question, Questioner Name, Questioner URL, Time of posting question, Question Category, Best Answer, Answerer Name, Answerer URL, Time of posting answer, and Rating of the answer. All the functions are written in the *contentscript.js* file.
2. Send data to the server. The server processes the data and returns the result. The client-side Javascript communicates with the server by sending a *XMLHTTPRequest*. The *POST* method is used to send the request because the data to be sent may be big for using the *GET* method. Also for data extracted from the *zhidao.baidu.com* website the encoding of the data is set to *gb2312* in order to encode Chinese characters.
3. The result is displayed to the user. If the user is an authorized user, the user can enter his/her feedback to the server.

5.3 Server Design

The server communicates with the plugin and also maintains a database system. The database system stores the information of Q&A sessions and the prediction model. The server receives the Q&A session data sent from the browser plugin. If the database has the label for the session, the server returns the label. If it is a new session, the server stores it in a buffer, calculates the spam grade based on the current model parameters, and returns the spam grade if necessary (i.e., a campaign session is detected). When enough data has been collected, we can use the *helpers* to label the data. Using logistic regression, the detection model will be updated using previous data as well as the newly labelled data.

5.4 Evaluation of Adaptive Online Detection System

To evaluate the performance of adaptive online detection system, we use the collected data from *Baidu Zhidao* and relay the data in iterations to simulate a real-world scenario. In particular, we pretend that initially we only have partial data and use the data as the training dataset to build a detection model. In each iteration, we add some new sessions and use them as the test dataset to test the performance of the detection system. At the end of an iteration, the new sessions are added into the training dataset, and the detection model is updated using the new training dataset. This step corresponds to the scenario that new data are labelled and added into the system. Then we repeat with another iteration.

For the test, we begin with 500-sample training set and build an initial detection model. At each iteration, we add 200-sample test set. After evaluating the detection performance, we expand the training dataset with the 200 test samples, and update the detection model with the new training dataset. We repeat this until we use up all 4998 samples.

Figure 8 and Figure 9 show the update of model parameters and the detection performance in each iteration. We




Figure 7: System architecture and communication between the client and the server

can observe that the detection model tends to converge after enough sessions have been added into the database over several iterations. This test scenario is similar to the practical application where we predicate the unknown sessions using current knowledge and train a new model based on the sessions after we manually give labels to them. The good performance results similar to those presented in Section 4.2 indicates that our system can effectively adapt model parameters to achieve good performance.




Figure 8: Adaptive changes of model parameters over time




Figure 9: The performance of the online detection system over time

To illustrate the advantage of adaptiveness, we also perform another test in which we fix the model after it is trained on the initial dataset. We use 200 samples as the initial training data and build a model. We fix the model parameters, and at each iteration, we test 200 new sessions using the fixed model. The results are shown in Figure 10

Compared with Figure 9, we observe nearly perfect precision but degraded performance on the rest of the metrics of the fixed model. We further looked into the values of




Figure 10: The performance of the fixed model

TP, TN, FP, FN and found that the false negatives were too high in the non-adaptive model. However, the false positives were very low. It meant that the non-adaptive model classified many campaign Q&A sessions as the non-campaign sessions. Consequently, although the precision is high, other decreased metrics implied that the non-adaptive model had obvious bias in classification. What's worse, this model cannot update itself by new samples because the parameters are only trained on the initial training dataset. Therefore, making the predication model adaptive to new samples is a necessary objective of the system design.

6. RELATED WORK

Our research is mostly related to work on spam detection and recognizing experts or authoritative users and trustworthy content in the social media. These topics have become crucial to many online services, especially the question and answer communities, whose contents are generated by millions of users. We discuss prior work on two aspects.

6.1 Retrieving high-quality answers in CQA sites

A lot of research has been done on finding high quality content in CQA sites. However, we haven't seen any paper which explicitly solved the credibility problem introduced in our work. Usually, researchers treated the best answers as the high-quality answers which has the risk of being defeated by the paid posters. In our work, we explicitly consider the credibility issues about the best answers.

Jeon *et al.* [13] attempted to predict the quality of answers in a community based question answering service with only non-textual features, such as *Answerer's Acceptance Ration*, *Answer Length* and *User's Recommendation*. They assumed the user feedback was a reliable source for the evaluation. Jurczyk *et al.* [15] presented a study of link structure of Yahoo! Answers [30]. They adopted an adaptation of the *HITS* algorithm [17] for finding experts in the Q&A portal. Their research was also based on the assumption that the

user feedback could be used to assign weights on the edges of their graph representing user relationships.

Liu *et al.* [18] applied their automated summary technique to summarize answers for questions which ask for opinions. They used cosine similarity to cluster topic-oriented answers and eliminated irrelevant ones. Bian *et al.* [6] tried to use both relevance between questions and answers and the quality of answers to retrieve good answers for a user query. Both textual features and statistical features such as user ratings were used in their approach. Later, in another work by Bian *et al.* [5], they explicitly considered the effect of several vote spam attacks. Such activities involved malicious voting for specific answers to improve their ranking and to decrease the ranking of competitors at the same time.

Agichtein *et al.* [2] studied the basic elements of social media and combined three features of the social media (Yahoo! Answers) to facilitate the task of identifying high quality content, namely intrinsic content quality, interactions between users and content usage statistics. *HITS* and *PageRank* algorithms were used to calculate the hubs and authorities users scores and usage statistics such as number of clicks of the Q&A session were used to complement the link-based analysis.

Fichman [9] conducted a comparative study of answer quality on multiple Q&A websites, Yahoo! Answers, Wiki Answers [27], Askville [3] and the Wikipedia Reference Desk [28]. Accuracy, completeness and verifiability were used as the quality measures for cross platform comparison. Fichman found that the quality of answers was significantly improved only in terms of answer completeness and verifiability, not the answer accuracy.

6.2 Other research work about crowd-sourcing spams but in different realms

Previous research has also investigated the crowd-sourcing spam in other areas. Jindal *et al.* [14], Ott *et al.* [21] and Arjun *et al.* [19] attempted to detect fake review or opinion spam in the online shopping stores, like Amazon's online store. Similar to research in CQA websites, they also used textual similarity features and user-oriented features, like ratings and history records. Huang *et al.* [12] developed a regression model with features suggesting quality-biased short text in Microblogging service, Twitter. They judged the quality of tweets based on relevance, informativeness, readability, and politeness of the short content and assigned different scores from 1 to 5. However, they didn't explicitly present how they define a spam-like tweet. Huang [11] conducted a similar study of commercial spam on blogging sites. They showed that the propaganda of some products in the comment of a blog post was crucial in detecting the malicious comments. The propaganda appeared in the form of URL, phone number, E-mail address, MSN numbers etc.

7. CONCLUSIONS AND FUTURE WORK

Detection of hidden campaigns can improve the user's experience when using current social websites. In this paper, we disclose the behavior of a specific group of online paid posters who create commercial campaigns on the community Q&A websites. We collect real-world datasets and identify effective features to distinguish normal sessions and the campaigns. The performance of our classifier, with integrated statistic and semantic analysis, is quite promising on the real-world case study. Based on a learning technique,

we also implement a prototype of adaptive online detection system which can retrieve the result in real time. The campaign scores and/or predicated labels can help users make better decisions when searching for answers on CQA portals and help the questioners select better answers as well.

This work is our first effort to detect online paid posters of CQA websites. In the future, we will test more features to improve the adaptive performance.

8. REFERENCES

- [1] *Applied Logistic Regression Analysis (Quantitative Applications in the Social Sciences) (v. 106)*. Sage Publications, Inc, 2nd edition.
- [2] E. Agichtein, C. Castillo, D. Donato, A. Gionis, and G. Mishne. Finding high-quality content in social media. In *Proceedings of the international conference on Web search and web data mining*, WSDM '08, pages 183–194, New York, NY, USA, 2008. ACM.
- [3] Askville. <http://askville.amazon.com/index.do>.
- [4] BaiduZhida. <http://zhidao.baidu.com/>.
- [5] J. Bian, Y. Liu, E. Agichtein, and H. Zha. A few bad votes too many?: towards robust ranking in social media. In *Proceedings of the 4th international workshop on Adversarial information retrieval on the web*, AIRWeb '08, pages 53–60, New York, NY, USA, 2008. ACM.
- [6] J. Bian, Y. Liu, E. Agichtein, and H. Zha. Finding the right facts in the crowd: factoid question answering over social media. In *Proceeding of the 17th international conference on World Wide Web*, WWW '08, pages 467–476. ACM, 2008.
- [7] J. Bian, Y. Liu, D. Zhou, E. Agichtein, and H. Zha. Learning to recognize reliable users and content in social media with coupled mutual reinforcement. In *Proceedings of the 18th international conference on World wide web*, WWW '09, pages 51–60, New York, NY, USA, 2009. ACM.
- [8] C. Chen, K. Wu, V. Srinivasan, and X. Zhang. Battling the internet water army: Detection of hidden paid posters. arXiv:1111.4297v1, 2011.
- [9] P. Fichman. A comparative assessment of answer quality on four question answering sites. *Journal of Information Science*, vol. 37 no. 5(476-486), 2011.
- [10] D. W. Hosmer and S. Lemeshow. *Applied logistic regression (Wiley Series in probability and statistics)*. Wiley-Interscience Publication, 2 edition, Sept. 2000.
- [11] C. Huang, Q. Jiang, and Y. Zhang. Detecting comment spam through content analysis. In *Proceedings of the 2010 international conference on Web-age information management*, WAIM'10, pages 222–233, Berlin, Heidelberg, 2010. Springer-Verlag.
- [12] M. Huang, Y. Yang, and X. Zhu. Quality-biased ranking of short texts in microblogging services. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 373–382, Chiang Mai, Thailand, November 2011. Asian Federation of Natural Language Processing.
- [13] J. Jeon, W. B. Croft, J. H. Lee, and S. Park. A framework to predict the quality of answers with non-textual features. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '06, pages 228–235, New York, NY, USA, 2006. ACM.
- [14] N. Jindal and B. Liu. Opinion spam and analysis. In *Proceedings of the international conference on Web search and web data mining*, WSDM '08, pages 219–230, New York, NY, USA, 2008. ACM.
- [15] P. Jurczyk and E. Agichtein. Discovering authorities in question answer communities by using link analysis. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, CIKM '07, pages 919–922, New York, NY, USA, 2007. ACM.
- [16] J. Kapur and H. Kesavan. *Entropy Optimization Principles with Applications*. Academic Press Inc., 1992.
- [17] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *J. ACM*, 46:604–632, September 1999.
- [18] Y. Liu, S. Li, Y. Cao, C.-Y. Lin, D. Han, and Y. Yu. Understanding and summarizing answers in community-based question answering services. In *Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1*, COLING '08, pages 497–504, Stroudsburg, PA, USA, 2008. Association for Computational Linguistics.
- [19] A. Mukherjee, B. Liu, and N. Glance. Spotting fake reviewer groups in consumer reviews. In *Proceedings of the 21st international conference on World Wide Web*, WWW '12, pages 191–200, New York, NY, USA, 2012. ACM.
- [20] Naver. <http://www.naver.com/>.
- [21] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, HLT '11, pages 309–319, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics.
- [22] M. S. Pera and Y.-K. Ng. A community question-answering refinement system. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*, HT '11, pages 251–260, New York, NY, USA, 2011. ACM.
- [23] Quora. <http://www.quora.com/>.
- [24] J. Surowiecki. *The Wisdom of Crowds*. Anchor, Aug. 2005.
- [25] Tiancaicheng. <http://www.tiancaicheng.com/>.
- [26] G. Wang, C. Wilson, X. Zhao, Y. Zhu, M. Mohanlal, H. Zheng, and B. Y. Zhao. Serf and turf: Crowdturfing for fun and profit. In *Proceedings of The 20th International World Wide Web Conference (WWW)*, 2012.
- [27] WikiAnswers. <http://wiki.answers.com/>.
- [28] WikipediaReferenceDesk. <http://en.wikipedia.org/wiki/wikipedia:referencedesk>.
- [29] Yahoo!Answers. 1 billion answers served!
- [30] Yahoo!Answers. <http://answers.yahoo.com/>.
- [31] Zhubajie. <http://www.zhubajie.com/>.